1365

Learning, Sticky Inflation, and the Sacrifice Ratio

John M. Roberts

June 2007

Kiel Institute for World Economics
Duesternbrooker Weg 120
24105 Kiel (Germany)

Kiel Working Paper No. 1365

Learning, Sticky Inflation, and the Sacrifice Ratio

by

John M. Roberts

June 2007

Learning, Sticky Inflation, and the Sacrifice Ratio

John M. Roberts[*]   Federal Reserve Board

April 23, 2007

Abstract

Over the past forty years, U.S. inflation has exhibited highly persistent movements. Moreover, these shifts in inflation have typically had real consequences, implying a "sacrifice ratio," whereby disinflations are typically associated with recessions and persistent increases in inflation often associated with booms. One hypothesis about the source of the sacrifice ratio is that inflation—and not just the price level—is sticky. Another is that private-sector agents typically must infer changes in inflation objectives indirectly from central bank interest-rate policy. The resulting learning process can lead to a sacrifice ratio trade-off. In this paper, I allow for both sticky inflation and learning in interpreting U.S. macroeconomic developments since 1955. Two key empirical findings are, first, that allowing for learning reduces the evidence for sticky inflation. Second, there is less evidence for sticky inflation in the post-1983 period than earlier. Indeed, in some estimates, there is little evidence of sticky inflation in the period since 1983, although this result is sensitive to the details of the specification. Nonetheless, simulation results suggest that for realistic models, the sacrifice ratio can be accounted for entirely by learning.

Since the 1960s, U.S. inflation has exhibited highly persistent movements. Barsky (1987) emphasized this phenomenon and it has been confirmed in recent research by Stock and Watson (2007) and Cogley and Sargent (2006). These low-frequency movements in inflation appear to have real consequences. As documented by Ball (1994), disinflationary episodes in the United States and other countries have typically been accompanied by output and employment losses. This correlation can be summarized by a sacrifice ratio, with the interpretations that a central bank which seeks to lower the average rate of inflation must be willing to accept a period of low output and employment.

As Taylor (1983) and Ball (1995) have pointed out, nominal rigidities, by themselves, are not sufficient to generate a sacrifice ratio. They show that if monetary policy is perfectly credible and transparent, disinflation can be carried out costlessly even when wages and prices are sticky. Indeed, Ball shows that under these conditions, it is theoretically possible for disinflation to cause a boom.

One feature of recent macroeconomic models that can give rise to a sacrifice ratio is sticky inflation. Sticky inflation has been added to macroeconomic models chiefly as a means of addressing some empirical shortcomings of sticky price models. For example, Fuhrer and Moore (1995) introduce sticky inflation so as to increase the predicted serial persistence of inflation. Various structural interpretations of sticky inflation have been, including real wage rigidities (Fuhrer and Moore, 1995; Blanchard and Gali, 2007), imperfect rationality (Roberts, 1997, 1998; Gali and Gertler, 1999; Mankiw and Reis, 2002), and indexation to past inflation (Christiano, Eichenbaum, and Evans, 2005; Smets and Wouters, 2003). While sticky inflation has been introduced primarily to address higher-frequency properties of inflation, Bomfim et al (1997) show that sticky inflation can give rise to a sacrifice ratio even when monetary policy is transparent and credible.

Another explanation for the real effects of low-frequency movements in inflation is learning: Ball (1995), Bomfim et al (1997), and Erceg and Levin (2003) have suggested that when the central bank's long-run inflation objective is unclear and agents must infer it from central bank actions, the resulting learning process can lead to output and employment losses when the central bank chooses to reduce inflation. The models of these papers assume sticky prices and as noted above, Bomfim et al (1997) also assume sticky inflation. But Ball (1995) and Erceg and Levin (2003) demonstrate that learning can lead to costly disinflation (a nonzero sacrifice ratio) even

when inflation isn't sticky.

To help further our understanding of the relative importance of sticky inflation and learning in accounting for the sacrifice ratio, this paper estimates a model with both features. In addition to sticky prices and (possibly) sticky inflation, the model incorporates other New Keynesian features including an IS curve with habit persistence and a monetary-policy reaction function.

Much recent empirical work emphasizes that the behavior of U.S. monetary policy and inflation have been very different before and after the early 1980s. Clarida, Gali, and Gertler (2000), for example, argue that U.S. monetary policy has been well-characterized by a Taylor rule since the early 1980s but fit that paradigm less well in earlier periods. Evidence provided by Stock and Watson (2007) and Cogley and Sargent (2006) suggests that U.S. inflation has become less persistent over this period as well. Because of these changes in policy and inflation dynamics, the empirical work emphasizes estimation over two periods, 1955 to 1983 and 1984 to 2004.

Two key findings are, first, that taking account of learning reduces the evidence in favor of sticky inflation. Second, there is less evidence for sticky inflation in the post-1983 period than earlier. Indeed, in some estimates, it appears that once learning is introduced, there is little evidence of sticky inflation in the post-1983 period. This result, however, is sensitive to the exact details of the specification.

The paper then evaluates the relative roles of learning and sticky inflation in accounting for the sacrifice ratio using model simulations. These simulations indicate that learning is the main source of costly disinflation: When there is learning, the sacrifice ratio is high even when we eliminate inflation stickiness from the model. But eliminating learning while preserving sticky inflation leads to a sharp drop in the sacrifice ratio.

The paper's penultimate section reconsiders the various explanations for sticky inflation in light of the finding that sticky inflation appears to have been an empirically important feature of the 1955-83 period but not more recently. Among the more-prominent explanations, indexation holds up best, because there is evidence that formal indexation has become less prevalent in the U.S. economy. That said, formal indexation was never very important; indexation can only fully account for sticky inflation if we are willing to accept that informal indexation was also widespread. Another possible explanation for the higher estimates of sticky inflation in the earlier period may be that because of the unstable policy environment, learning may have been more difficult than predicted by the model. Some preliminary work comparing

3

model forecasts with surveys of inflation expectations suggests that this may indeed have been the case.

This paper has points in common with several other recent papers. Milani (2006) also estimates a model with learning and sticky inflation. However, Milani looks at a general learning process and does not specify the aspect of the economy about which agents must learn. Here, by contrast, agents learn about a specific quantity—the central bank's inflation objective. Another related paper is Erceg and Levin (2003). They argue that learning can account for the sacrifice ratio during the Volcker disinflation of 1980-83. They do so in a calibrated model, however, and they do not discuss how their model would perform in other periods. Ireland (2006) also examines a New Keynesian model with a time-varying inflation target and learning. But Ireland does not draw out the implications of his model for the sacrifice ratio.

1 The model

The model involves three observable variables—the output gap y, inflation Δp, and the short-term interest rate r. These variables are linked through the following New Keynesian-style model:

$$y_t = \frac{1}{1+\eta}E_t y_{t+1} + \frac{\eta}{1+\eta}y_{t-1} \tag{1}$$
$$-0.25\frac{1-\eta}{1+\eta}\sigma(r_{t-1} - E_{t-1}\Delta p_t - rstar) + u_t$$

$$u_t = \phi_1 u_{t-1} + \varepsilon_t^y \tag{2}$$

$$r_t = \rho r_{t-1} + (1-\rho)\{rstar + (\Delta p_t + \Delta p_{t-1} + \Delta p_{t-2} + \Delta p_{t-3})/4$$
$$+\lambda_\pi[(\Delta p_t + \Delta p_{t-1} + \Delta p_{t-2} + \Delta p_{t-3})/4 - pitarg_t] \tag{3}$$
$$+\lambda_y y_t + \lambda_{dy}(y_t - y_{t-1})\} + a_1\Delta r_{t-1} + a_2\Delta r_{t-2} + \varepsilon_t^r$$

$$\Delta p_t = \frac{\beta}{1+\beta\omega}E_t\Delta p_{t+1} + \frac{\omega}{1+\beta\omega}\Delta p_{t-1} + \frac{(1-\beta)(1-\omega)}{1+\beta\omega}pibar \tag{4}$$
$$+4\frac{(1-\zeta)(1-\beta\zeta)}{\zeta(1+\beta\omega)}\nu y_t + \varepsilon_t^p$$

$$pitarg_t = \phi_2 pitarg_{t-1} + (1-\phi_2)pibar + \varepsilon_t^{targ} \tag{5}$$

Equation 1 is a New Keynesian IS curve. Although it is an equation for the output gap, its microeconomic foundations are those of the Euler equa-

tion for consumer spending with (external) habit persistence. As discussed in Woodford (2003), because consumer spending is the largest single component of spending in the United States, overall spending appears to be well approximated by a model for consumer spending. Under that interpretation, the parameter η is the degree of habit persistence and the parameter σ is related to the curvature of the utility function. The interest rate term is premultiplied by 0.25 because interest rates and inflation are annualized, but Euler equations are typically estimated on non-annualized data. The IS-curve error term is allowed to be serially correlated (equation 2).

Equation 3 is the monetary-policy reaction function. It is similar to equations estimated by others, such as Clarida, Gali, and Gertler (2000), English, Nelson, and Sack (2003), and Gorodnichenko and Shapiro (2006). Notably, it includes lags of the funds rate as well as the change in the output gap. Also, it assumes that policy reacts only to current and lagged values of output and inflation. As Gorodnichenko and Shapiro emphasize, with these features, the policy rule nests both the familiar Taylor rule and the price-level-targeting rules advocated by, among others, Woodford (2003). The results of Gorodnichenko and Shapiro suggest that recent U.S. monetary policy is well-characterized as a weighted average of these policies. There is a residual shock to the reaction function, reflecting movements in monetary policy not predicted by the explicit arguments of the function. In addition, the central bank's inflation objective, *pitarg*, may also be subject to shocks. We will return to the question of what process *pitarg* may follow shortly.

Equation 4 extends the Calvo model of price determination to allow for partial indexation of prices that are not re-optimized to past inflation, as in Smets and Wouters (2003). The parameter ω measures the degree of indexation; when $\omega = 1$, there is complete indexation, as in Christiano, Eichenbaum, and Evans (2005) and when $\omega = 0$, there is no indexation to past inflation, as in the original Calvo model. The parameter ζ is related to the frequency of price adjustment, with the average interval between price adjustments equal to $\frac{1}{1-\zeta}$. The parameter ν relates marginal cost to the output gap. This parameter will be affected by the slope of labor supply, but also by such factors as the demand elasticity of individual firms and the degree to which capital and other fixed factors are firm-specific; see Eichenbaum and Fisher (2004) for a discussion. In equation 4, ζ and ν are not separately identified. I will therefore assume throughout that $\zeta = 0.68$, which corresponds to an average interval of price adjustment of about three quarters. Note that when $\omega = 0$, there is still indexation, to the long-run average inflation rate

pibar. As noted by Yun (1996), such indexation is needed in the Calvo model to prevent too wide a dispersion of prices across firms when there is ongoing inflation. Note, however, that because the discount factor β is close to one, the impact of this steady-state indexation on inflation dynamics is minimal.[1]

I consider several specifications for the evolution of the central bank's inflation objective, *pitarg*. These are nested in equation 5. One possibility is that $\phi_2 = 0$; in this case, the inflation target is a constant. Another possibility is that *pitarg* follows a random walk ($\phi_2 = 1$). This assumption has received a great deal of attention in recent empirical work; it is the assumption made, for example, by Stock and Watson (2007), Ireland (2006), and Cogley and Sargent (2006). Strictly speaking, however, the assumption of a random walk process for *pitarg* is not compatible with equation 4, because there is no well-defined long-run inflation rate (*pibar*) in this case. I consider two ways of resolving this tension. One is to preserve the random walk assumption and eliminate the role of *pibar* in equation 1 by assuming $\beta = 1$. The other is to compromise on the random walk assumption by assuming that ϕ_2 is close to, but slightly less than, one. As we will see in section 3, unfortunately, these competing assumptions affect the empirical results.[2]

When the inflation target is time-varying, the model can capture persistent, low-frequency movements in inflation. But because the model includes shocks to the reaction function, there will be learning in this case: When agents see an unexpected change in the federal funds rate, they do not know whether it is the result of a transitory shock to policy (an ε^r shock) or a shift in the inflation target (an ε^{targ} shock). Agents, like we econometricians, are assumed to use the Kalman filter to come up with their estimates of the proportion of each shock the reaction-function surprise represents.

In this model, the central bank's inflation objective has an important

---

[1]When consumer preferences are marked by habit persistence, marginal cost will be affected by changes in consumer spending as well as by the level, and should reflect restrictions imposed by the estimated coefficients of the consumer spending Euler equation. I do not impose these restrictions, however, because while equation 1 is motivated by the consumer problem, it is in fact estimated with output data and, as noted by Woodford (2003, ch. 5), the coefficients of an output Euler equation can be very different from those of consumption, even when the Euler equation broadly captures output dynamics.

[2]Ireland (2006) takes an alternative approach to resolving this tension. He eliminates any role for *pibar* in the Calvo model by allowing indexation to *pitarg*. One problem with this assumption, however, is that while there is some evidence that wages and prices have been indexed to lagged inflation, there is little evidence that they have been indexed to time-varying estimates of underlying inflation.

influence on inflation through the general equilibrium solution to the model: Current inflation is affected by expected future inflation, which in turn is affected by expected future output gaps. Expectations of future output gaps will be affected by monetary policy—and, in particular, by the relation of expected inflation to the (perceived) inflation target. It is only when inflation is expected to line up with the target that output gaps will converge to zero. Agents will thus recognize that the path of future output gaps will be such as to ensure that inflation will eventually converge to the inflation target.

2 Estimation with a fixed inflation target

In this section, I consider estimation of the model with a fixed inflation target. As noted in the introduction, because inflation has been highly persistent, it is unlikely that a fixed inflation target is a good characterization of the data. However, it provides a useful benchmark, as it is relatively straightforward to implement and has been used in other studies.

Table 1 presents estimates of the model with a fixed inflation target over three periods: 1955-2004, 1955-1983, and 1984-2004. In the estimation, inflation is measured as the annualized quarterly percent change in the price index for personal consumption expenditures, the federal funds rate is the short-term interest rate, and the output gap is estimated by running the Hodrick-Prescott filter ($\lambda = 16,000$) through the log of real GDP; the result is multiplied by 100 to make the units comparable to inflation and the interest rate. The discount factor β in equation 4 is assumed to be 0.99, a standard assumption. The model was estimated with full-information maximum likelihood, using the FIML option of the Dynare program.

The first column of table 1 presents results over the full sample. The IS-curve slope parameter σ was not precisely estimated, and I imposed a value of 0.3. This value is consistent with micro and macro evidence on the intertemporal elasticity of substitution in consumer spending—see, for example, Elmendorf (1996) and Elmendorf and Mankiw (1998).[3] The estimated habit persistence parameter η is 0.89,suggesting a high degree of habit persistence. There is also a moderate degree of serial persistence in the residuals—$\phi = 0.34$.[4]

---

[3]Unconstrained, the estimate of σ would be 0.2 and the log likelihood would be -810.4. The other estimated parameters are little affected.

[4]For the IS curve, there is an alternative local maximum. At this alternative local

Table 1: Model Estimates with Fixed Inflation Target

	(1)	(2)	(3)	(4)	(5)
	1955-2004	1955-83	1984-2004	1955-83	1984-2004
η	.89 (.04)	.85 (.06)	.96 (.03)	.86 (.07)	.95 (.03)
σ	.30 (–)	.30 (–)	.30 (–)	.30 (–)	.30 (–)
ϕ	.34 (.09)	.33 (.12)	.23 (.12)	.33 (.12)	.26 (.12)
ω	.92 (.05)	.98 (.08)	.61 (.10)	.96 (.05)	.61 (.10)
ν	.025 (.017)	.034 (.028)	.016 (.015)	.025 (–)	.025 (–)
λ_y	1.45 (.43)	2.12 (1.05)	.65 (.32)	2.13 (1.02)	.67 (.32)
λ_{dy}	1.96 (.84)	1.72 (1.30)	2.23 (1.15)	1.75 (1.31)	2.15 (1.11)
λ_{dp}	.42 (.23)	.46 (.35)	.68 (.45)	.47 (.36)	.67 (.44)
ρ	.89 (.02)	.90 (.04)	.90 (.03)	.90 (.04)	.89 (.03)
sd_y	.31 (.04)	.39 (.06)	.19 (.03)	.39 (.06)	.19 (.03)
sd_{dp}	.66 (.03)	.71 (.05)	.62 (.06)	.71 (.05)	.63 (.06)
sd_{ff}	.80 (.04)	.96 (.06)	.42 (.03)	.96 (.06)	.42 (.03)
log-L	-810.5	-523.6	-223.7	-523.7	-223.9

Taken literally, the estimate of ω implies that 92 percent of prices are indexed to past inflation each period and there is thus a very high degree of inflation stickiness. The estimate of the elasticity of marginal cost with respect to the output gap, $\nu = 0.025$, is not very precise. It is also fairly small. By way of comparison, as reported in Woodford (2003), Rotemberg and Woodford (1997) found an estimate of $\nu = 0.14$, considerably larger than the estimate reported in column 1.

The estimates of the policy reaction function imply a high degree of interest-rate smoothing ($\rho = 0.89$), strong reaction to the level of the output gap ($\lambda_y = 1.5$), and a moderately high degree of responsiveness to inflation ($\lambda_{dp} = 0.4$). There is also a fairly strong reaction to the change in the level

maximum, habit persistence is relatively low and the degree of serial persistence in the residuals is relatively high. Empirically, it is difficult to distingusih these two local maximums: In some sample periods and specifcations, high η and low ϕ are preferred. In all cases, however, the likelihoods are close. In principle, these parameters can be distinguished because they have different implications for the effect of interest rates on output. But as noted in the text, the IS slope parameter σ is not precisely estimated, making it difficult to distinguish the two local maximums. Here, I have chosen to focus on one local maximum across all samples to facilitate comparison of results. The estimates of the parameters of the Phillips curve and reaction function are little affected by the different IS curve results.

of economic activity, with $\lambda_{dy} = 2.0$.

Column 2 presents results for the 1955-83 period. The estimated habit persistence and serial correlation in IS-curve shocks are similar to those in column 1. The degree of indexation ω is close to one and the elasticity of marginal cost with respect to output (ν) is 0.034, larger than in column 1 but less precisely estimated. The reaction-function coefficients are similar to those in column 1, with considerable interest-rate smoothing, large coefficients on both the level and the change in economic activity, and a modest response to inflation.

Column 3 presents results for the 1984-2004 period. The IS-curve estimates shift a bit from those in the first two columns; the estimated degree of habit persistence is higher while the degree of serial correlation of the residuals is a bit lower. The differences in the inflation parameters are sharper. In particular, the share of indexing ω drops from nearly one in the early sample to 0.6 in the post-1983 sample; it remains strongly statistically significant. In addition, the inflation-equation parameter ν drops by about one-half relative to the estimate in column 2. For the reaction function, the main change is in the coefficient on the level of the output gap, where λ_y drops by two-thirds, to 0.7. The estimated weight on the change in the output gap rises somewhat, as does the coefficient on inflation. [5]

As might be expected given the "Great Moderation" in the economy's volatility, the standard deviation of the shock to the IS curve falls by about one-half after 1983. The standard deviation of the shock to the reaction function also falls, by more than one-half. However, the standard deviation of the shock to the inflation equation is only about 12 percent smaller in the later sample.

Columns 4 and 5 of the table examine the sources of change in the parameters of the inflation equation more closely. In the 1984-2004 period, the point estimates of both ω and ν fell. Given the lack of strong theoretical underpinnings for sticky inflation, the drop in ω is perhaps not too surprising.

---

[5]The similarity of the estimated policy reaction functions across the two periods is surprising, given the results of Clarida, Gali, and Gerlter (2000), who find much smaller estimates of the coefficients on output and inflation in the earlier period. However, the model specification and sample examined here are different from those of CGG. It is noteworthy that these differences had little impact on the post-1983 estimates, which are similar to those of CGG. That estimates for the early period are sensitive to details of specification is consistent with the notion discussed in section 6, that it is difficult to characterize policy in the early period.

Our confidence in a structural interpretation of the New Keynesian model would be enhanced, however, if the slope parameter ν were the same in the two periods. Although the point estimates of ν are different in the two samples, the estimates are not very precise, suggesting that the differences may not be important statistically. In contrast, the difference in the point estimates of ω is large relative to their estimated standard errors. In columns 4 and 5, ν is constrained to equal its full-sample value, 0.025, which is between the estimates for the two subsamples. As can be seen, with this restriction imposed, the likelihood in either sample barely changes, suggesting that a constant ν is not at variance with the data.

3 Adding a time-varying inflation target

This section presents estimates of the model with a time-varying inflation target. Here, the main parameters of the reaction function are constrained to equal their values in the post-1983 sample.[6] Under this assumption, monetary policy in the 1960s and 1970s was conducted under the same general principles as in the post-1983 period. So, all of the difference in policy across the two periods results from variation in the shocks to policy, both to the implicit inflation target and elsewhere. This perspective on shifts in policy differs from that of Clarida, Gali, and Gertler (2000), who argue that the parameters of the reaction function changed. It is closer, however, to the view of Orphanides (2001). Orphanides argues that policy followed the same basic principles in both periods but that mismeasurement of potential output in the early period accounted for the poor macroeconomic performance of that period. Here, we assume that it is mostly the behavior of the shocks to monetary policy.

The model in equations 1 through 5 includes three observables but four shocks. To estimate the model (as well as the implicit inflation target), the Kalman filter is used. As before, the estimation is performed using the FIML options of the Dynare program. As noted in section 1, in estimating the model with the Kalman filter, we are implicitly assuming that private agents also use the Kalman filter in forming their estimate of the central bank's

---

[6]The exception is for the coefficients on the change in the federal funds rate. These differ considerably across the two periods, suggesting important high-frequency differences in how policy was implemented. These differences do not have an important effect on the lower-frequency properties of policy that are our concern here.

Table 2: Model Estimates with Random-Walk Inflation Target

	(1)	(2)	(3)	(4)
	1955-83	1984-2004	1984-2004	1984-2004
Inflation measure:	Overall	Overall	Overall	Core
η	.83 (.06)	.96 (.04)	.95 (.04)	.94 (.04)
σ	.30 (–)	.30 (–)	.30 (–)	.30 (–)
ϕ	.31 (.11)	.23 (.12)	.24 (.13)	.26 (.13)
ω	.81 (.11)	.05 (.15)	.30 (.14)	.06 (.17)
ν	.029 (.023)	.031 (.022)	.024 (.018)	.019 (.014)
sd_y	.38 (.06)	.19 (.03)	.19 (.03)	.19 (.03)
sd_{dp}	.73 (.05)	.83 (.11)	.76 (.09)	.64 (.10)
sd_{ff}	1.01 (.07)	.43 (.03)	.43 (.03)	.43 (.03)
sd_{targ}	.36 (.13)	.24 (.08)	.08 (–)	.08 (–)
log-L	-525.4	-226.5	-230.0	-206.1

inflation target. Because the Kalman filter is the best available method for discerning the inflation target, private agents can thus be considered to be learning optimally.

3.1 Inflation target follows a random walk

Table 2 presents results for the case of $\phi_2 = 1$ and $\beta = 1$; we will turn shortly to the case of $\phi_2 = 0.999$ and $\beta = 0.99$. For the IS curve, the results for the 1955-83 period shown in column 1 are similar to those in table 1: Once again, there is a high degree of habit persistence and a moderate degree of serial correlation. The key new feature introduced here—time-variation in the inflation target, as captured by a non-zero value for the standard deviation of the inflation target, sd_{targ}—is strongly statistically significant.[7] The estimate of ν is somewhat smaller than in table 1, although, again, this parameter is not precisely estimated. The estimated indexation share is 0.8. It remains highly statistically significant. But this estimate is almost two standard deviations less the estimated value in table 1. The reduction in the estimate of ω suggests that optimal learning can account for a portion of the estimated degree of inflation stickiness in this period.

---

[7]A comparison of the likelihoods in tables 1 and 2 would appear to suggest that allowing for a time-varying inflation target has reduced the fit of the model somewhat. But recall that the other coefficients of the reaction function are constrained in table 2.

11

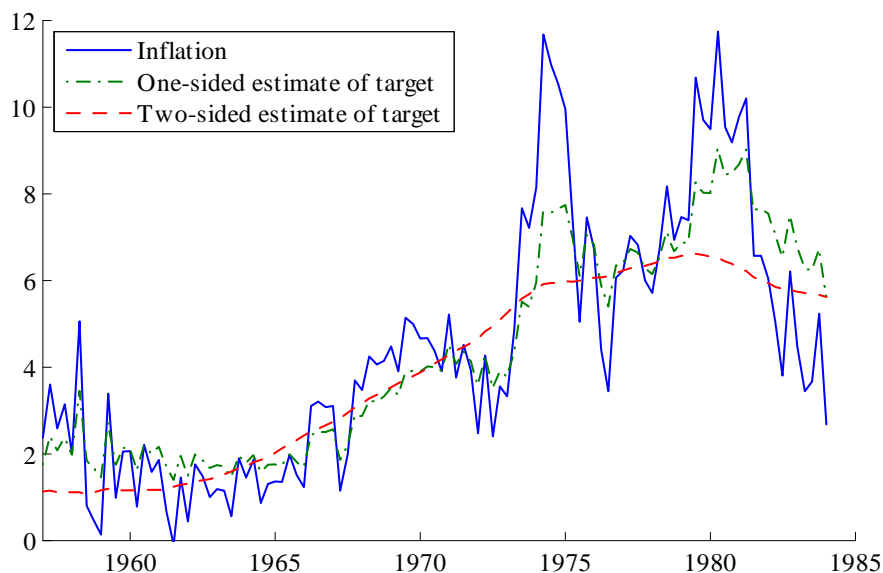Figure 1: Inflation and Estimated Target Inflation, 1957-1983

Figure 1 shows the one-sided and two-sided estimates of the inflation target implicit in the estimates in column 2. The two-sided estimate of the target rises from around 1 percent in the late 1950s and early 1960s to a bit more than 6 percent in the second half of the 1970s. In the early 1980s, the two-sided estimate moves down, edging below 6 percent by the end of 1983. The real-time (one-sided) estimates are more variable and indicate that this model ascribes considerable variation to the Fed's inflation target. Nonetheless, it is worth noting that for the period when survey estimates of long-run inflation expectations begin to become available—starting in 1981—there is broad agreement between these estimates: Like the one-sided estimates, the long-run expectations of professional forecasters also move down from around 8 percent in 1981 to 6 percent by 1983 (see discussion in section 5). It thus appears that, at least by the early 1980s, optimal filtering in a model like this one leads to estimates of the implicit inflation target that are consistent with the survey evidence.[8]

---

[8]Erceg and Levin (2003) report a similar finding.

Column 2 shows results for the post-1983 sample. Looking first at the shocks, the standard deviations of both monetary-policy shocks are smaller than in the earlier period. However, the decline in the standard deviation of the white-noise shock is proportionally larger than for the inflation target. The IS-curve parameters are similar to those in column 4 of table 1.

The estimates of inflation dynamics in column 2 are striking: Here, the degree of indexation is estimated to be close to zero. The slope coefficient is higher than in column 3 of table 1 and is similar to that in column 1. These results suggest that in the post-1983 sample, allowing for optimal learning entirely removes the evidence of a moderate degree of indexation that was found when we assumed a fixed inflation target, as in table 1.

Figure 2 presents estimates of the implicit inflation target consistent with

Figure 2: Inflation and Estimated Target Inflation, 1985-2004
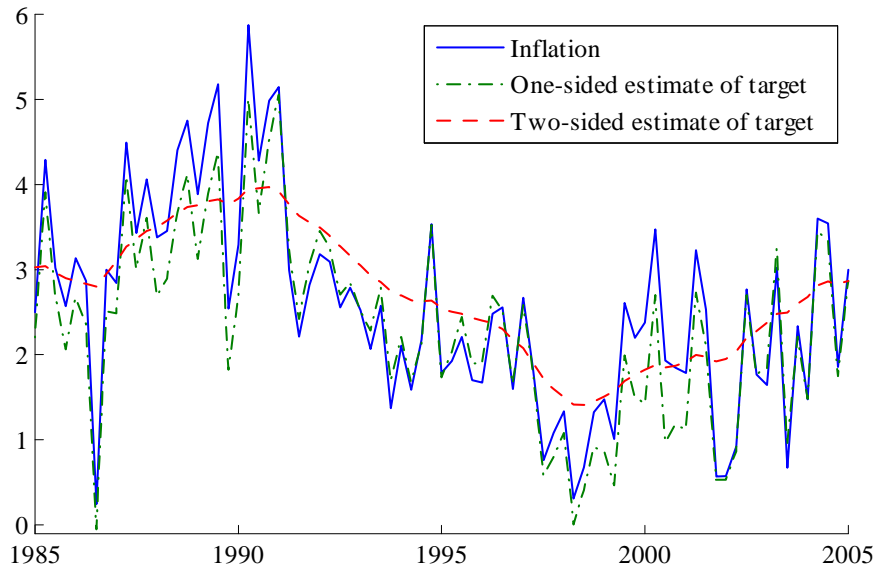
the estimates in column 3 of table 2. As can be seen, the implicit target is quite variable, with the smoothed (two-sided) estimate varying between $1\frac{1}{2}$ and $3\frac{1}{2}$ percent over the post-1983 sample; the one-sided estimate is nearly as variable as inflation itself. While there is widespread agreement that target

13

inflation varied over this period—and, in particular, stepped down around the early 1990s—the movements in figure 2 seem far too large.

To address the high variability of the one-sided estimate of the inflation target, in column 3, the standard deviation of the target inflation shock is constrained to be 0.08. The fit of the model deteriorates in this case; evidently, greater variation in the implicit inflation target is preferred. Most of coefficient estimates are not affected by this restriction. The exception is the indexation share, which is now estimated to be 30 percent; it is statistically significant at conventional levels. Nonetheless, this is substantially less indexation than was estimated with a fixed inflation target.

In column 4, core inflation (excluding food and energy) is used in place of overall PCE inflation. It is of interest to explore core inflation in the current context because one possible source of the high estimated indexation parameter may be serial persistence in energy-price shocks, in particular, to crude-oil prices. As in column 3, the standard deviation of inflation-target shocks is constrained to be 0.08. As can be seen, the estimated indexation parameter is once again quite small, suggesting that serial persistence in energy-price shocks may indeed have led to a spurious finding of significant indexation in column 3. The estimate of ν is also smaller than in column 3, possibly because energy prices are more cyclically sensitive than other consumer prices. Another notable change from the estimates in column 3 is in standard deviation of the shock to the price equation, which is smaller, reflecting the lower volatility of core inflation. (As a consequence, the estimated log likelihood is also smaller; it should thus not be compared with the log likelihood in columns 2 and 3.)

Figure 3 presents estimates of the inflation target consistent with the estimates in column 4 of table 2. The two-sided estimate of the inflation target moves down from around $3\frac{3}{4}$ percent in the late 1980s to around $2\frac{1}{4}$ percent for the period since 1995, similar to other estimates of trend inflation, such as those of Levin and Piger (2004). There continues to be considerable variation in the one-sided estimate of the target, however.

3.2 Inflation target highly persistent, but ultimately mean-reverting

In this subsection, I turn to the case in which $\beta = 0.99$ and $\phi_2 = 0.999$. Thus, the inflation target is highly persistent but ultimately mean-reverting.

Figure 3: Core Inflation and Estimated Trend, Reduced Variance, 1985-2004
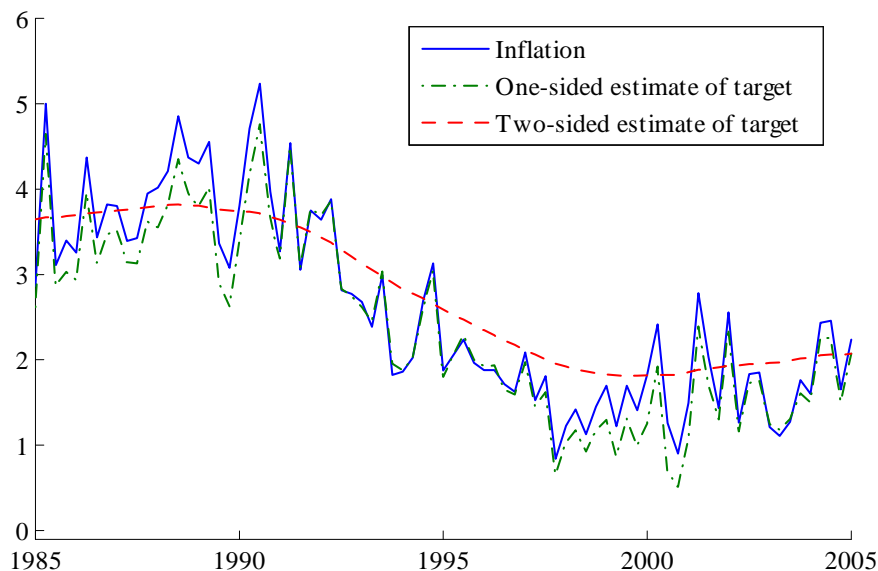
These assumptions allow logic of the Calvo model to be preserved, with some cost to the assumption that the inflation target follows a random walk that has been made in other recent work.

Table 3 presents results. For the 1955-83 period (column 1), the results are very similar to those in table 2: The degree of indexation ω is slightly larger than 0.8 and the elasticity of marginal cost with respect to the output gap ν is around 0.03. For the 1984-2004 period, however, the results are somewhat different than before. In particular, in column 2, the point estimate of ω is now equal to 0.23. While the t-ratio is only 1.3, this estimate is much larger than in table 2. Moreover, in column 3, when the standard deviation of the shock to the inflation target is restricted to be 0.08, ω is estimated to be around 0.5 and strongly statistically significant.[9]

While the estimates in tables 2 and 3 suggest that there is some sensitivity of the results to the exact details of the specification, overall, the results suggest that allowing for optimal learning reduces the evidence for sticky

---

[9]Estimates with core inflation were similar to those in column 3.

Table 3: Model Estimates with Highly Persistent, but Stationary, Inflation Target

	(1)	(2)	(3)
	1955-83	1984-2004	1984-2004
η	.83 (.06)	.93 (.04)	.94 (.05)
σ	.30 (−)	.30 (−)	.30 (−)
ϕ	.31 (.11)	.27 (.13)	.25 (.15)
ω	.83 (.10)	.24 (.18)	.52 (.11)
ν	.033 (.021)	.035 (.019)	.022 (.016)
sd_y	.38 (.06)	.19 (.03)	.19 (.04)
sd_{dp}	.73 (.05)	.70 (.09)	.65 (.18)
sd_{ff}	1.01 (.07)	.43 (.03)	.42 (.04)
sd_{targ}	.34 (.12)	.30 (.13)	.08 (−)
log-L	-528.5	-223.0	-225.7

inflation. This is especially true in the recent period. In the pre-1984 period, however, there remains considerable evidence of sticky inflation even after taking account of optimal learning about a changing inflation target.

4   Learning and the sacrifice ratio

In the model estimated in the previous section, there were permanent changes in the central bank's inflation objectives that were not immediately evident to agents. As noted by Ball (1995), Bomfim et al (1997), and Erceg and Levin (2003), in this case, disinflation can be costly. That's because, in this model, the central bank can only signal its policy intentions through changes in interest rates. Because such shocks to policy are only sometimes related to changes in the inflation target, it can take a while for agents to sort out whether any given shock to policy is the result of a persistent shock to the inflation target or owes to some other source. During this transition period, agents put some weight on the possibility that movements in interest rates are the result of transitory shocks to monetary policy, which, as in most New Keynesian models, have effects on real economic activity.

One way of characterizing the costs associated with disinflation is the sacrifice ratio—that is, the cumulative output gains or losses associated with a permanent rise or fall (respectively) in the inflation target. The sacrifice

ratio is a reduced-form quantity that will be affected by many parameters in the model. Of particular interest here are the speed of learning and the degree of inflation stickiness, which, as noted in the introduction, can both lead to costly disinflation. In the extreme case of immediate recognition of a change in the inflation target, no inflation stickiness, and no lags in the monetary-policy rule, the sacrifice ratio will be zero.

To calculate the sacrifice ratio, we need to know how quickly the public's perceptions of the inflation target respond to changes in the central bank's target, based on their filtering of residuals to the reaction function. The most efficient way to do so is to apply the Kalman filter. In particular, equations 3 and 5 imply that the reaction-function error will be:

$$u_t = \varepsilon_t^r + (1 - \rho)\lambda_\pi \widehat{\varepsilon}_t^{targ} \tag{6}$$

where $\widehat{\varepsilon}_t^{targ}$ is the inflation-target forecast error. Assuming that the inflation target follows a random walk ($\phi_2 = 1$ in equation 5), the public should apply the following formula so as to update their estimate of the inflation target according to the Kalman filter:

$$\widehat{pitarg}_t = \widehat{pitarg}_{t-1} + \frac{\gamma}{(1 - \rho)\lambda_\pi} u_t \tag{7}$$

where γ is the (steady-state) Kalman gain,

$$\gamma = \frac{\phi}{2}\left(\sqrt{1 + \frac{4}{\phi}} - 1\right) \tag{8}$$

and ϕ is the signal-to-noise ratio,

$$\phi = \left[\frac{(1 - \rho)\lambda_\pi \sigma_{targ}}{\sigma_r}\right]^2 \tag{9}$$

Table 4 shows the effects of a 1 percentage point change in target inflation for various values of the Kalman gain. Because the sacrifice ratio is defined with respect to permanent changes in the inflation target, it is appropriate to use parameter estimates from table 2. I focus on the estimates for the 1955-83 period (column 1) because most estimates of the sacrifice ratio are based on this period (see Ball, 1994, for example). Those parameters imply a Kalman gain of $\gamma = 0.025$. As can be seen in the second line of the table, agents learn about the target very slowly in this case—even twenty years after the initial shock, the perceived inflation target is still only 0.86

17

Table 4: Effects of a 1 percentage change in the inflation target under learning. Fixed indexation ($\omega = 0.8$), varying gain.

Kalman gain (γ)	Expected target Years since shock			Inflation Years since shock			Sacrifice ratio Years since shock		
	5	10	20	5	10	20	5	10	20
.0125	.23	.39	.63	.36	.50	.69	2.2	3.8	4.9
.025	.41	.64	.86	.52	.72	.89	1.7	2.8	3.4
.05	.65	.87	.98	.77	.92	.99	1.4	2.2	2.5
.10	.89	.99	1.0	1.0	1.0	1.0	1.3	1.9	2.0
1.0	1.0	1.0	1.0	1.1	1.0	1.0	1.3	1.5	1.5

percent. Because of this slow convergence, the sacrifice ratio will depend on the horizon at which it is measured: After ten years, the sacrifice ratio indicates that 2.8 percentage points of annual lost output were associated with each percentage point reduction in inflation; after twenty years, the sacrifice ratio is 3.4. These estimates of the sacrifice ratio are in the ballpark of conventional estimates for this period: Typical estimates of an output-based sacrifice ratio in the pre-1984 period run from 3 to 5 (employment-based output ratios are smaller—around 2—reflecting the well-known Okun's law phenomenon).[10]

In the first row of the table, the Kalman gain is half the estimated value. This case might be of interest if agents believed the inflation target to be less variable than was in fact the case. As will be discussed in section 6, the 1960s and 1970s were a period during which monetary policy was difficult to characterize. As a consequence, learning in that period may have been less than optimal. With this learning speed, the convergence of the perceived inflation target is very slow, as is that of actual inflation. Estimates of the sacrifice ratio are on the high end of the conventional estimates, in the range of 4 to 5 at a ten-to-twenty-year horizon.

The value of the standard deviation of the shock to the inflation target for the 1955-83 period—0.36—is on the low side of other estimates of the variability of permanent shocks to inflation for this period. Stock and Watson (2007), for example, find estimates ranging from 0.4 to more than 1.0 over this period. The third and fourth lines of the table shows the implications of higher standard deviations of the inflation target, of 0.50 and 0.72, respectively. With the larger gains in these cases, agents learn about the shift in

---

[10]Bomfim et al (1997) review sacrifice-ratio estimates.

the inflation target faster, and the sacrifice ratio is smaller.

The Kalman gain of 0.1 in line 4 is similar to the value estimated by Erceg and Levin (2003) over the 1981-to-1985 period. With this gain, learning is virtually complete in ten years and the sacrifice ratio is only 1.9. Erceg and Levin report an estimated sacrifice ratio of 1.7 over this period. Ball (1994) calculates a similar sacrifice ratio, of 1.8, over the 1980-83 period. Ball also argues that the sacrifice ratio was particularly low in this period and suggests that faster learning may have lowered the sacrifice ratio.

The estimates for the post-1983 period in table 2 indicated a drop in the volatility of both the inflation target and the other shocks to monetary policy. In column 2 of table 2, the drop in the estimated volatility of the other shocks to monetary policy was proportionally greater than that of the drop in the shock to the inflation target. By themselves, these changes in volatility estimates would imply an increase in the Kalman gain—to 0.038— and thus a drop in the sacrifice ratio—to around $2\frac{1}{2}$ at a ten-year horizon. As noted in the previous section, however, the implicit inflation target in this case was implausibly variable. Under the alternative assumptions in column 3 of table 2, the Kalman gain would be 0.0125, implying a somewhat higher sacrifice ratio than the baseline estimates.

The final row of table 3 shows the implications of immediate recognition of changes in the inflation target. From the perspective of the model estimated here, immediate recognition cannot be achieved, because it is impossible to distinguish the two shocks to monetary policy in real time. Still, this case provides a useful benchmark for comparison. With immediate recognition, actual inflation converges to target within five years. The sacrifice ratio is 1.5 in this case.

Table 5: Effects of a 1percent increase in inflation target under learning. Fixed gain ($\gamma = 0.025$), varying indexation.

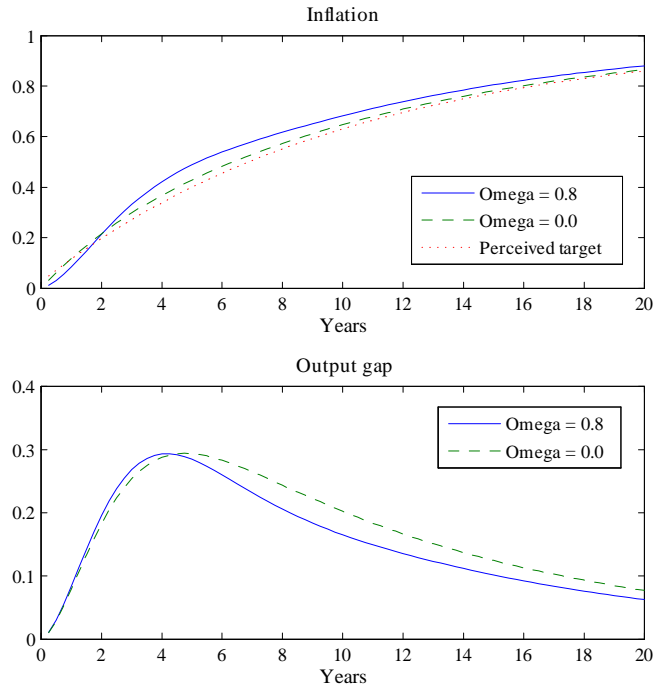|	Inflation			Sacrifice ratio		
	Years since shock			Years since shock		
ω	5	10	20	5	10	20
0.0	.45	.67	.87	2.0	3.3	4.1
0.5	.48	.69	.88	1.9	3.0	3.8
0.8	.52	.72	.89	1.7	2.8	3.4
1.0	.54	.76	.90	1.8	2.6	3.0

Table 5 presents the implications of variation in the degree of indexation, holding the Kalman gain fixed at its estimated value of 0.025. (Because the evolution of the target is little different from that in the second line of table 3, it is omitted here.) As discussed in the introduction, indexation is in principle an alternative source of costly disinflation. But in this model, varying the degree of indexation has little effect on the sacrifice ratio: With no indexation, the ten-year sacrifice ratio is 3.3, actually a bit larger than with the baseline indexation of 0.8. This finding is similar to that of Erceg and Levin (2003), who also found that they could duplicate the empirical sacrifice ratio in a model without sticky inflation.

Figure 4 illustrates why indexation makes so little difference to the sacrifice ratio. The solid lines depict the evolution of inflation and output following a 1 percentage point drop in the inflation target under the baseline estimates from column 1 of table 2—in particular, with the Kalman gain equal to 0.025 and $\omega = 0.8$. The dashed lines show the effects with ω set equal to zero. As can be seen, the main features of the simulations are similar: In each case, inflation moves slowly toward its long-run objective and there are large output losses associated with the transition. As might be expected, inflation moves very closely with the perceived target in the case of no inflation stickiness. By contrast, when $\omega = 0.8$, inflation initially moves a bit more sluggishly than the target and then overshoots. However, the main dynamics are determined by the sluggishness of learning, and these high-frequency differences do not affect the sacrifice ratio very much.[11]

The results presented in this section suggest that learning has likely been more important than sticky inflation in accounting for the costs of disinflation: Starting from realistic model parameters, reducing indexation to zero actually boosted the sacrifice ratio somewhat. By contrast, moving from the estimated pace of learning to immediate recognition of a change in the

---

[11]The failure of indexation to have much effect on the sacrifice ratio may come as a surprise. Appendix A explores this issue for a broader range of model parameters. Those simulations suggest that in the case of a monetary policy without lags and with immediate recognition of shifts in the inflation target, the degree of indexation has the expected effect on the sacrifice ratio, with the sacrifice ratio rising from zero in the case of no indexation to notable levels with the indexation parameter in the range of 0.8 to 1.0. However, the estimates presented in sections 2 and 3 indicate that lags in monetary policy are important, and that changes in inflation targets may have been difficult to discern from changes in interest rates in the 1955-83 period. As the simulations in tables 4 and 5 suggest, in these more-realistic settings, the sacrifice ratio is affected to a greater degree by learning than by the degree of indexation.

Figure 4: Effects of a Permanent Increase in the Inflation Target, Gain = 0.025

inflation target reduced the long-term sacrifice ratio by a factor of two or more.

5   Moving-average expectations

In section 3, agents were assumed to estimate the inflation target using optimal learning as represented by the Kalman filter. While optimal learning could, in some specifications, account for the observed degree of inflation stickiness in the post-1983 period, it could not in the earlier period. One possibility is that, in a period such as 1955-83, when monetary policy was difficult to understand, agents used other rules for forming their long-run inflation expectations.

One conjecture about how agents form their long-run inflation expecta-

tions is that they use moving averages of past inflation. Kozicki and Tinsley (2001) have argued that a weighted average of past inflation with geometrically declining weights does a good job of matching survey measures of long-run inflation expectations. And Stock and Watson (2007) argue that an IMA(1,1) model is a very good univariate model of inflation over the 1953-to-2004 period. An implication of the IMA(1,1) model is that the long-run inflation target is equal to a geometrically weighted moving average of past inflation, so the Kozicki-Tinsley and Stock-Watson characterizations of long-run inflation expectations are very similar. There is, however, an important difference: Over the 1955-83 period, the estimates of Stock and Watson imply that the moving-average weights drop off at a very steep rate—at least 50 percent per quarter and sometimes as high as 90 percent, effectively making their forecast of long-run inflation equal to last period's inflation rate. By contrast, Kozicki and Tinsley suggest a much shallower rate of decline in the weights, of about $1\frac{1}{2}$ percent per quarter.

Figure 5 illustrates the implications of different weighting schemes for the estimate of long-run inflation along with long-run inflation expectations from surveys of professional forecasters.[12] As can be seen, depending on the period, the SPF has been well-approximated by a geometric moving average with weights that decline at a pace of either 5 or 10 percent per quarter. When the weights decline at a rate of 20 percent per quarter, the implicit inflation target follows actual inflation more closely and thus does not match the SPF very well. Hence, long-run expectations based on very short moving averages, such as those proposed by Stock and Watson for the early 1980s, do not line up well with the available survey evidence.

Table 6 shows estimates of the model over the 1955-83 period with a geometric moving average of past inflation serving the role of *pitarg* in equation 3. Three measures of the weighted average of past inflation are used, with weights that decline at rates of 5, 10, and 20percent per quarter. Because target inflation implicitly has a unit root in this case, I again impose the restriction $\beta = 1.0$, as in table 2. When the weights decline at a pace of 5 percent per quarter, the estimated degree of inflation stickiness is about

---

[12]This series reports results for five-to-ten-year forecasts of consumer-price inflation. For the period from 1981 to 1991, economist Richard Hoey conducted a survey of forecasters for his brokerage firm. For the period since 1991, the series is from the Philadelphia Fed's Survey of Professional Forecasters. The surveys collected CPI forecasts; in recent years, the CPI has increased on average about 1/2 percentage point more than the PCE price index used in this paper, and the plotted series adjusts for this bias.

Figure 5: Professional Forecasters' vs. Geometric Moving-Average Estimates of Inflation Target
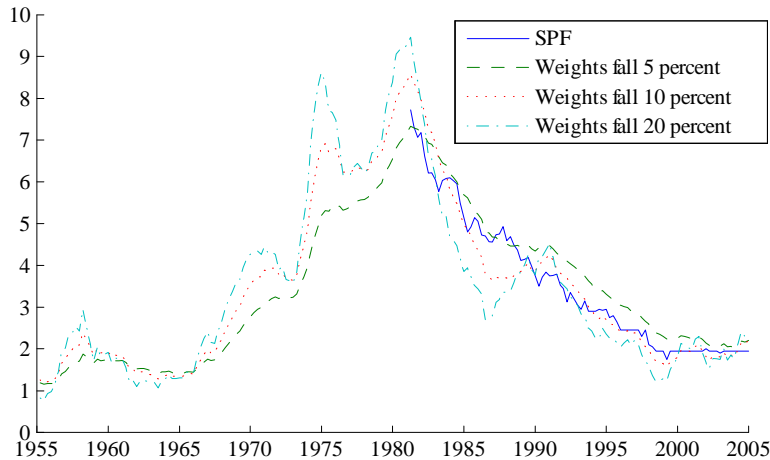
80 percent, comparable to that for the learning model reported in table 2. The main slope parameter of the price equation—ν—is estimated to be larger than before and is now statistically significant at conventional levels. Otherwise, parameter estimates are similar to those in table 2. As the rate of decline in the moving-average weights increases, the estimated degree of inflation stickiness falls, so that, when the weights decline 20 percent per quarter, the indexation weight is estimated to be 50 percent. The model fit, as indicated by the log-likelihood, is highest in this case. Again, the slope of the price equation is large and statistically significant.

The bottom row of the table shows the sacrifice ratio implied by each model. When the moving-average weights decline slowly, the sacrifice ratios are higher than the estimate (of three) implied by the baseline learning estimates in column 1 of table 2. However, at least when the weights decline at 10 percent per quarter, the resulting sacrifice ratio is in line with the conventional wisdom. When the weights decline at 20 percent per quarter, the sacrifice ratio is similar to that associated with the baseline learning model.

Overall, the results in table 5 suggest that when we limit ourselves to moving-average estimates that do a good job of following empirical estimates of long-run inflation expectations, this model does not help reduce the

Table 6: Model estimates with moving-average inflation target
1955-83

	(1)	(2)	(3)
Moving average parameter:	.05	.10	.20
η	.82 (.06)	.83 (.06)	.83 (.06)
σ	.30 (–)	.30 (–)	.30 (–)
ϕ	.33 (.12)	.32 (.12)	.32 (.12)
ω	.83 (.07)	.69 (.07)	.48 (.08)
ν	.047 (.022)	.052 (.020)	.052 (.017)
sd_y	.38 (.06)	.38 (.06)	.39 (.06)
sd_{dp}	.71 (.05)	.71 (.05)	.70 (.06)
sd_{ff}	1.01 (.07)	1.01 (.07)	1.02 (.07)
log L	-518.2	-510.2	-495.9
10-year sacrifice ratio	5.9	4.0	3.0

estimated degree of inflation stickiness very much. However, the resulting sacrifice ratios are plausible, and the model fits somewhat better. When we allow the implicit long-run inflation target to follow actual inflation closely, the moving-average inflation-target model yields a smaller degree of indexation.

6 Why Is Inflation Less Sticky?

6.1 Positive models of sticky inflation

As noted in the introduction, there are a number of hypotheses about the sources of sticky inflation. What does the reduction in inflation stickiness documented in sections 2 and 3 suggest about the plausibility of these various stories? The decline in inflation stickiness would seem to argue against explanations that rely primarily on inherent structural properties of the economy. Of the proposed explanations for inflation stickiness, the most structural are the real-wage-rigidity explanations emphasized by Fuhrer and Moore (1995) and more recently by Blanchard and Gali (2007). There is no particular reason to expect real-wage rigidity to vary over time and, in particular, it shouldn't be related to changes in average inflation or monetary policy.

Nonrationality in the formation of expectations is less inherently structural, and could be affected by changes in policy. One way of examining

directly the degree of rationality in expectations formation is to look at survey measures of inflation expectations. Roberts (1998) and Carroll (2003) find that inflation expectations as captured by the Michigan survey deviate considerably from perfect rationality. Roberts argues that these deviations from perfect rationality are consistent with the observed degree of inflation stickiness. Carroll has examined how the degree of imperfect rationality in the Michigan survey has changed over time. He finds that as inflation has become low and stable, inflation expectations of households have become less rational, in the sense that they have deviated more from those of professional forecasters. Such a movement over time in the degree of rationality seems inconsistent with a decline in the degree of inflation stickiness.

Indexation has been proposed as another explanation for inflation stickiness. In particular, in the model proposed by Smets and Wouters (2003), firms may, in each period, either change prices optimally, move prices with lagged inflation (indexing), or leave prices unchanged in nominal terms. The available data suggests that indexation has become less prevalent in the U.S. economy. In particular, the degree of indexation in union wage agreements fell from about 60 percent in the late 1970s and early 1980s to around 20 percent in the mid-1990s, when the data ceased to be collected. An important caveat, however, is that the fraction of workers covered by formal collective bargaining agreements was never more than 25 percent of the U.S. workforce. Evidence presented in Blinder et al (1998) suggests that, at least by the 1990s, formal indexation was rare outside the unionized sector of the economy.

Neither real rigidities nor imperfect rationality seem consistent with the observed change in the degree of inflation stickiness. Still, either may provide an explanation for some underlying level of sticky inflation, which some estimates suggest remained important in the post-1983 sample. Each of these explanations has independent support. As noted above, Roberts (1998) and Carroll (2003) argue that survey expectations suggest that expectations formation may be imperfectly rational. And Hall (2005) has argued that real-wage rigidity can help improve the empirical plausibility of labor-search models. Changes in the degree of indexation in the U.S. economy certainly go in the right direction to be consistent with some reduction in inflation stickiness. But formal indexation was never widespread in the U.S. An appeal to indexation as the main explanation for the drop in inflation stickiness would have to rely on considerable informal indexation in the 1960s and 1970s.

6.2 Policy confusion

The formal econometric modeling in section 3 was based on the premise that, aside from variation in the inflation target, economic policy was otherwise stable. But as Romer and Romer (2002, 2004) have argued, economic policy in the 1961-1981 period was anything but stable. At the start of that period, the Kennedy Administration explicitly advocated a New Economic Policy. That policy explicitly advocated using aggregate demand stimulus to achieve lower unemployment; considerations of the inflationary consequences were of secondary importance.

By the late 1960s, high inflation was a concern of policymakers, suggesting that the NEP proved to be more inflationary than expected. This inflationary surprise led central bankers and other policymakers to question their understanding of how the economy functions. As Romer and Romer document, among the aspects of the inflation process that were called into question were the natural rate of unemployment, the sacrifice ratio, and the role of supply shocks, in addition to the well-known debate over a trade-off between inflation and unemployment. In particular, over much of this period, policymakers were too optimistic about the levels of the natural rate and potential output (see Orphanides, 2001, for a discussion). Policymakers were also very pessimistic about the sacrifice ratio (Romer and Romer, 2002). Given their pessimism about the sacrifice ratio, the effects of supply shocks on inflation were thought to be intractable.

One implication of the policy instability of this period is that there may have been much for private-sector agents to learn about. In the model in section 1, all of these policy shifts can be summarized by shifts in the central bank's inflation objective. As the preceding brief review of the historical record suggests, that's probably an oversimplification. One possible reason for the appearance of sticky inflation in this period is that agents were not able to keep pace with the variations in policy—and thus made systematic, and serially correlated forecast errors. Because these forecast errors were serially correlated, the model in section 1 may have interpreted them as evidence of sticky inflation.

Figure 6 provides one bit of evidence on whether forecast errors can account for the large estimate of ω in the 1955-83 period. Figure 6 compares forecast errors from the Livingston survey with forecast errors from two versions of the model. The model was used to generate one-year-ahead forecasts of inflation, assuming the parameter estimates presented in table 2 and the

Figure 6: Forecast Errors, Final Quarter of Each Year
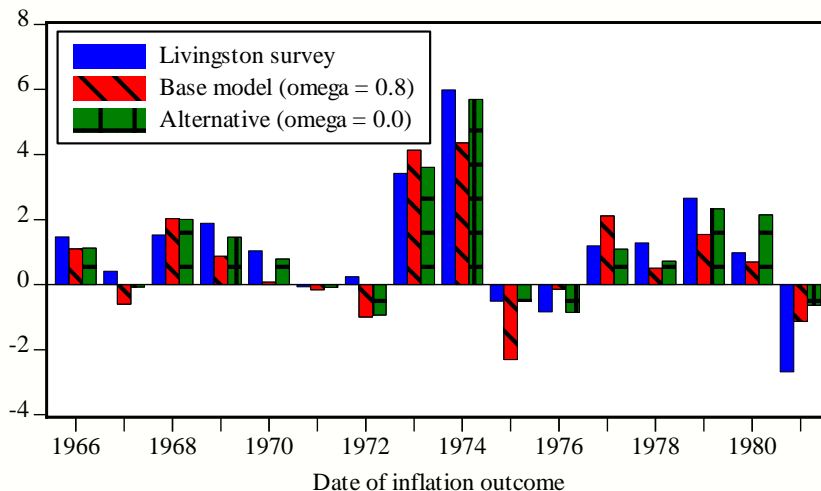
one-sided estimates of the inflation target presented in figure 1.[13] Forecasts were generated both with and without the estimated degree of indexation— that is, with $\omega = 0.8$ and $\omega = 0.0$. Figure 6 presents forecast errors over the 1966-to-1981 period.[14]

The model forecast errors and those from the Livingston survey are quite similar. In particular, inflation was systematically higher than expected throughout this period, regardless of the forecast method: The average forecast error for the Livingston survey was 1.2 percentage points; for the baseline model, it was a bit lower, at 0.9 percentage point, while for the alternative model with $\omega = 0.0$, it is also 1.2 percentage points. Moreover, the RMSE for the Livingston survey forecasts is actually a bit smaller than for the model forecasts. The large errors for the model over this period suggest that there is nothing especially "non-rational" about the Livingston survey forecast errors of this period, since even a model with rational expectations—and the

---

[13]The stationary part of model was solved in terms of data available at period t. This model solution was then used to forecast the deviation of inflation from its trend. This forecast was then combined with the period t estimate of the inflation target to obtain an inflation forecast.

[14]For clarity, figure 6 only shows forecast errors for the final quarter of each year. Results based on mid-year forecasts are similar.

27

Table 7: Summary Statistics for Forecast Errors
Forecasts Made 1965:H1 to 1981:H2

Forecast:	Mean error	RMSE
Livingston survey	1.17	1.98
Base model ($\omega = 0.8$)	.87	2.03
Alternative ($\omega = 0.0$)	1.18	2.12

benefit of in-sample fit—made similar forecast errors. One interpretation of this result is that the simple metaphor of an unknown, time-varying inflation target appears to generate considerable "policy confusion."

Examining the results more closely, there are a number of instances in which the alternative model does a better job of matching the Livingston survey's forecast error than does the baseline model; this may account for the better match on the mean error for the alternative model. A simple regression of the Livingston survey forecasts on the two model forecasts gives a substantial edge to the alternative model: The coefficient on the alternative model forecast is 1.09, while that on the baseline forecast is only 0.02. Because the model with $\omega = 0.0$ matches the Livingston survey more closely, these results provide some evidence that forecast errors can account for the large value of ω estimated in the pre-1984 period.

7   Conclusion

Sticky inflation and a time-varying inflation target are competing explanations for inflation persistence. Allowing for either can, in principle, lead to important costs associated with disinflation—that is, a sacrifice ratio. When there is a time-varying inflation target, it may be difficult for private-sector agents to infer it from central bank actions. The resulting learning process—even if it reflects optimal filtering—can lead to a sacrifice ratio. Alternatively, inflation stickiness—such as that associated with indexation—means that inflation cannot move promptly when the central bank's inflation target changes, even when the inflation target is known. The estimates of this paper suggest that learning has been much more important in accounting for the sacrifice ratio than has sticky inflation. In particular, the estimated degree of learning can account for conventional estimates of the sacrifice ratio even when inflation is not sticky. In contrast, if there is no learning, even high degrees of inflation stickiness cannot account for the observed sacrifice ratio.

The empirical work suggests that there has been some variation over time in the degree of inflation stickiness. In particular, inflation appears to be less sticky in the post-1983 period than was the case earlier. Indeed, some estimates suggest that inflation stickiness may have disappeared altogether in the most recent period—although this result is sensitive to the details of the specification. Nonetheless, the drop in the estimated degree of inflation stickiness suggests that the economic mechanism underlying sticky inflation is sensitive to the economic environment. As was discussed in section 6, that makes some explanations for sticky inflation more plausible than others. Of the leading explanations, indexation seems most consistent with the drop in inflation stickiness, as formal indexation has also fallen over time. Confusion about the objectives of policy in the 1960s and 1970s, beyond the confusion about the central bank's inflation objectives that lead to learning in this model, may be another reason for the estimated degree of inflation stickiness in the pre-1984 period.

One useful extension of this work would be to consider alternative models of sticky prices. In particular, the results for the post-1983 sample proved sensitive to assumptions made about long-run indexation. In the Calvo model, long-run indexation is crucial, because prices can remain fixed for extended periods. In other models of inflation—such as the staggered-contracts models of Taylor (1983) or Wolman (2004)—long-run indexation is less important, because prices cannot remain fixed as long as in the Calvo model.

As noted in section 4, the learning model of this paper has the formal implication that increases in the signal-to-noise ratio might lower the costs of disinflation—and Ball (1994), among others, has argued that such a mechanism was at work in the early 1980s. But the estimates of Stock and Watson (2007) suggest that the reason for the increase in the signal-to-noise ratio at that time was a sharp increase in the volatility of the permanent component of inflation. Presumably, raising the volatility of the inflation target would not be an attractive way of reducing the sacrifice ratio. A more promising way of increasing the signal-to-noise ratio of monetary-policy actions might be for a central bank to improve its communications. The results of this paper have nothing to say about the efficacy of such polices. The experience of inflation targeters in this regard is perhaps more germane: Despite a high degree of emphasis on communicating their long-run inflation objectives, countries adopting inflation targets have not found disinflation to be costless, or even notably less costly than might otherwise have been expected (Bernanke et al, 1999).

A Further Results on the Sensitivity of the Sacrifice Ratio

Table A.1 provides additional perspective on the sensitivity of the sacrifice ratio to the Kalman gain and the degree of indexing in the inflation equation. As in text tables 4 and 5, the parameter estimates for the IS curve and the inflation equation are taken from column 1 of table 2. (The second row and third column thus reproduce earlier results.) The sacrifice ratio is quite sensitive to the gain for each of the values of ω, although the sensitivity is smaller for larger values of ω. For intermediate values of the Kalman gain, the sacrifice ratio is fairly insensitive to the degree of indexation. For very small values of the gain, the sacrifice ratio falls with the degree of indexation, whereas under immediate recognition—a gain of 1.0—the sacrifice ratio is increasing in indexation, albeit only slightly.

Table A.1: Ten-year Sacrifice Ratio with Empirical Reaction Function

Kalman gain	ω			
	.0	.5	.8	1.0
.0125	5.0	4.4	3.8	3.2
.025	3.3	3.0	2.8	2.6
.05	2.4	2.3	2.2	2.4
.10	2.0	1.9	1.9	2.1
1.0	1.5	1.5	1.5	1.9

It may be surprising that the sacrifice ratio is not zero in the lower left-hand corner of table A.1, when the gain is one and there is no indexation. Here, an immediately recognized change in the inflation target continues to have implications for real output owing to the lags in the policy reaction function. In particular, the estimated monetary policy rule includes several lags of the federal funds rate and assumes that monetary policy reacts to a trailing four-quarter moving average of inflation. Table A.2 repeats the sensitivity analysis presented in table A.1 under alternative assumptions about monetary policy: The lagged funds rate is omitted from the reaction function and monetary policy is assumed to react to contemporaneous inflation only; otherwise, the parameters of the policy reaction function are the same as in tables 4 and 5.

	ω			
Kalman gain	.0	.5	.8	1.0
.0125	1.69	1.67	1.71	2.18
.025	.87	.88	.98	1.60
.05	.46	.47	.60	1.28
.10	.24	.26	.41	1.14
1.0	.00	.06	.23	1.01

Table A.2: Ten-year Sacrifice Ratio with Simplified Reaction Function

In table A.2, when there is immediate recognition of changes in the target rate of inflation and no indexation, the sacrifice ratio is zero. Both indexation and slow learning lead to positive sacrifice ratios. It is still the case, however, that moving to empirically relevant estimates of the gain has a larger impact on the sacrifice ratio than does an increase in ω from 0.0 to 0.8.

The sacrifice ratios in table A.2 are considerably smaller than those in table A.1. For example, with $\omega = 0.8$ and a Kalman gain of 0.025, the sacrifice ratio is 1.0 in table A.2 but 2.8 in table A.1. These results indicate that lags in monetary policy—such as those in the reaction function estimated here—tend to boost the sacrifice ratio, other factors equal.

References

[1] Ball, Laurence (1994) "What Determines the Sacrifice Ratio?" in *Monetary Policy*, N. Gregory Mankiw, ed., University of Chicago Press, 155-194.

[2] _____ (1995) "Disinflation with Imperfect Credibility," *Journal of Monetary Economics* 35, 5-23.

[3] Barsky, Robert B. (1987) "The Fisher Hypothesis and the Forecastability and Persistence of Inflation," *Journal of Monetary Economics* 19, 3-24.

[4] Bernanke, Ben, Thomas Laubach, Frederic Mishkin, and Adam Posen (1999). *Inflation Targeting: Lessons from the International Experience.* Princeton University Press.

[5] Blanchard, Olivier, and Jordi Gali (2005) "Real Wage Rigidities and the New Keynesian Model," working paper presented at Federal Reserve Board conference, *Quantitative Evidence on Price Determination*, September 29-30, 2005.

[6] Blinder, Alan, Elie Canetti, David E. Lebow, and Jeremy B. Rudd (1998) *Asking about Prices*, Russell Sage Foundation: New York.

[7] Bomfim, Antulio, Robert J. Tetlow, Peter von zur Muehlen, and John C. Williams (1997) "Expectations, Learning, and the Cost of Disinflation," Federal Reserve Board FEDS working paper no. 1997-42 (August).

[8] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113, 1-45.

[9] Clarida, Richard, Jordi Gali, and Mark Gertler (2000) "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115, 147-80.

[10] Cogley, Timothy, and Thomas J. Sargent (2006) "Inflation-Gap Persistence in the United States," unpublished manuscript (September).

[11] Cogley, Timothy, and Argia M. Sbordone (2006) "Trend Inflation and Inflation Persistence in the New Keynesian Phillips Curve," unpublished manuscript, Federal Reserve Bank of New York (July).

[12] Eichenbaum, Martin, and Jonas D. M. Fisher (2004) "Evaluating the Calvo Model of Sticky Prices," NBER working paper no. 10617 (June).

[13] Elmendorf, Douglas W. (1996) "The Effect of Interest-Rate Changes on Household Saving and Consumption: A Survey," Federal Reserve Board FEDS working paper no. 1996-27.

[14] Elmendorf, Douglas W., and N. Gregory Mankiw (1999) "Government Debt," in *Handbook of Macroeconomics*, volume 1C, John B. Taylor and Michael Woodford, eds., Elsevier Science, 1615-69.

[15] English, William B., William R. Nelson, and Brian P. Sack (2003) "Interpreting the Significance of Lagged Interest Rates in Estimated Monetary Policy Rules," *Contributions to Macroeconomics* 3, issue 1.

[16] Erceg, Christopher J., and Andrew T. Levin (2003) "Imperfect Credibility and Inflation Persistence," *Journal of Monetary Economics* 50, 915-44.

[17] Fuhrer, Jeff, and George Moore (1995) "Inflation Persistence," *Quarterly Journal of Economics* 110, 127-59.

[18] Gali, Jordi, and Mark Gertler (1999) "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44, 195-222.

[19] Gorodnichenko, Yuriy, and Matthew D. Shapiro (2006) "Monetary Policy When Potential Output Is Uncertain: Understanding the Growth Gamble of the 1990s," NBER working paper no. 12268 (May).

[20] Hall, Robert J. (2005) "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review* 95, 50-65.

[21] Ireland, Peter N. (2006) "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," NBER working paper no. 12492 (August).

[22] Kozicki, Sharon, and Peter A. Tinsley (2001) "Shifting Endpoints in the Term Structure of Interest Rates," *Journal of Monetary Economics* 47, 613-52.

[23] Levin, Andrew T., and Jeremy M. Piger (2004) "Is Inflation Persistence Intrinsic in Industrial Economies?" European Central Bank working paper no. 334 (April).

[24] Mankiw, N. Gregory, and Ricardo Reis (2002) "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117, 1295-1328.

[25] Milani, Fabio (2006) "A Bayesian DSGE Model with Infinite-Horizon Learning: Do 'Mechanical' Sources of Persistence Become Superfluous?" *International Journal of Central Banking* 2, 87-106.

[26] Orphanides, Athanasios (2001) "Monetary Policy Rules Based on Real-Time Data," *American Economic Review* 91, 964-85.

[27] Roberts, John M. (1997) "Is Inflation Sticky?" *Journal of Monetary Economics* 39, 173-96.

[28] Roberts, John M. (1998) "Expectations and the Transmission of Monetary Policy," Federal Reserve Board FEDS working paper no. 1998-43 (November).

[29] Romer, Christina D., and David H. Romer (2002) "The Evolution of Economic Understanding and Postwar Stabilization Policy," in *Rethinking Stabilization Policy*, Federal Reserve Bank of Kansas City, 11-78.

[30] _____ and _____ (2004) "Choosing the Federal Reserve Chair: Lessons from History," *Journal of Economic Perspectives* 18, 129-62.

[31] Sargent, Thomas J. (1999) *The Conquest of American Inflation.* Princeton University Press.

[32] Smets, Frank, and Raf Wouters (2003) "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association* 1, 1123-75.

[33] Stock, James H. and Mark W. Watson (2007) "Has Inflation Become Harder to Forecast?" *Journal of Money, Credit, and Banking* forthcoming.

[34] Taylor, John B. (1983) "Union Wage Settlements during a Disinflation," *American Economic Review* 73, 981-93.

[35] Wolman, Alexander L. (2001) "A Primer on Optimal Monetary Policy with Staggered Price Setting," *Federal Reserve Bank of Richmond Economic Quarterly* 87, 27-52.

[36] Woodford, Michael (2003) *Interest and Prices*, Princeton University Press.

[37] Yun, Tack (1996) "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics* 37, 345-370.